intel

AP-477

APPLICATION NOTE

Low Voltage Embedded Design

JOHN WILLIAMS
APPLICATIONS ENGINEER

February 1993

Order Number: 272324-001

Contact your local Intel sales office or your distributor to obtain the latest specifications before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained from:

Intel Corporation
P.O. Box 7641
Mt. Prospect, IL 60056-7641

or call 1-800-879-4683

Low Voltage Embedded Design

CONTENTS

CONTENTS

1.0 INTRODUCTION

Currently, integrated circuit manufacturers are specifying devices for low voltage operation. Most devices are specified for operation around 3V. This can indicate operation centered around 3.0V, 3.3V or other voltages. Many devices are even specified for operation at 2V and below.

In many embedded designs, the designer wants the lowest system power dissipation possible. Embedded, low voltage designs require maximum battery life which is directly proportional to the current consumption of the system. System current consumption is a linear function of operating voltage.

This application note discusses why low voltage devices and systems are beneficial and how to design low voltage systems. Power supplies, mixed voltage systems, single voltage systems and power management schemes are all discussed. Although 3V systems are specifically discussed, the concepts apply to all low voltage designs.

1.1 What is 3V?

The term "3V" usually refers to 3.3V ±0.3V. This is the JEDEC standard for regulated supplies. The majority of devices are specified for 3.3V operation, the standard being used for portable PC's. JEDEC Standard 8 defines LVCMOS (Low Voltage CMOS) operating voltages and interface levels for low voltages. JEDEC Standard 8.1 defines LVTTL (Low Voltage TTL) compatibility. The LVTTL standard defines specifications for low voltage devices that operate in 5V TTL systems without interface logic.

There is also a standard for unregulated supplies, 2.7V to 3.6V (battery operated systems). Few devices are currently specified for operation in the unregulated region. Many devices used in embedded applications have a wider V_{CC} range and do not conform to this standard. A V_{CC} specification from 2.7V to 5.5V allows operation from an unregulated alkaline battery supply (3 x AA). Older 5V devices derated for 3V operation (specified for 3V operation at a reduced frequency) typically can operate in this region.

2.0 REASONS FOR 3V OPERATION

Components are being specified to operate at 3V for two reasons. First, advanced fabrication technologies incorporating smaller device geometries require lower operating voltages. Second, there are a number of advantages to operating a system at low voltage.

2.1 Device Geometries

One of the driving factors to reduce component operating voltages is technology. As device channel lengths and oxide thicknesses decrease (Figure 2-1), lower operating voltages are required to maintain component reliability. Three main reliability concerns are: oxide breakdown, punch-through, and hot-electron effects.

As feature length and width in a component decrease, depth also decreases. With thinner gate oxides, lower gate voltages are required to avoid dielectric breakdown in the oxide. If oxide breakdown occurs, the gate of the device begins to draw current.

As the channel lengths shorten, lower source and drain voltages are required to keep the depletion regions around the wells from meeting, at this point, the current through the device is no longer controlled by the voltage on the gate.

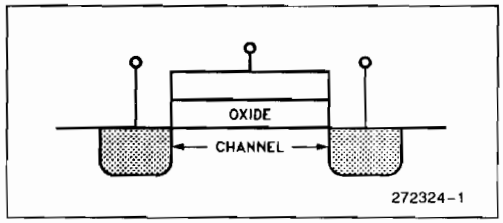

Figure 2-1. Basic Transistor Structure

Finally, lower voltages are required to avoid hot-electron effects. Electrons can lodge in the gate oxide, altering the charge in the oxide. When charge is added to the gate oxide, the turn-on voltage of the device is altered, degrading performance.

2.2 System Benefits

Presently, not all devices are being designed to operate specifically at low voltages. Many devices are 5V designs with their performance derated for low voltage operation. These devices are only rated for low voltage operation because of the benefits it has for system designs. Operating parts at lower voltages has some significant advantages at a system level.

Battery powered systems operate significantly longer at 3V. Battery life, rated in Amp-hours, is a function of the current a system draws. The current consumption of a system varies linearly with the operating voltages of the devices. Equation 2.1 expresses this relationship.

Equation 2.1: $I_{CC} = C \int V \bullet dt \cong \Delta V \bullet C \bullet F$

where

I_{CC} = Device current consumption

ΔV = Switching voltage

C = Device capacitance and output capacitive loading

F = Device operating frequency

Determining the current consumption of an entire system is not quite this straightforward, but this topic will be examined in more detail later. The point is, lowering device operating voltages from 5V to 3V decreases current consumption by at least 40 percent. Battery life increases by more than 40 percent, as batteries last longer at lower discharge rates. For example, a battery rated at 100 mAH with a load of 10 mA would last 10 hours. If the battery only had to supply 5 mA, it would have a life longer than the expected 20 hours. Reducing system voltage also decreases the number of batteries needed to produce the required operating voltage.

Heat generated by a device is proportional to the power it consumes. The formula for power consumption of a device is the current (Equation 1) multiplied by the operating voltage (V). Power consumption varies as the square of the operating voltage. Reducing the operating voltage of a device from 5V to 3V decreases power consumption by 64 percent. Heat dissipation is reduced by the same amount. This has a couple of important implications. On a device level, parts that generate less heat do not have as many constraints placed on packaging. Packages with much smaller footprints can be created and plastic packages can be offered for devices previously only offered in ceramic packages. On a system level, if parts produce less heat, they can be placed closer together. These advantages allow reduction of the overall form factor and weight of a system. Additionally, many device failure mechanisms are heat related. Systems producing less heat have higher reliability.

Noise generated by devices is a concern for all system designers. The effects of noise: overshoot, undershoot, ground bounce, etc., are a function of dV/dt. On low voltage devices, dV/dt is lower. This is discussed in more detail in a later section. Decreased noise generation by low voltage systems simplifies design and makes meeting maximum noise limits easier.

The move to lower operating voltages has a number of benefits to the system designer. Smaller, lighter, less noisy systems and longer battery life all contribute to better designs. Low voltage designs also enable older designs to become mobile/portable.

3.0 LOW VOLTAGE SYSTEM DESIGN

Low voltage systems have many aspects in common with standard 5V designs. In addition to the normal design considerations made in any 5V system, low voltage designs have some areas that need specific attention. System power supply, mixed voltage interfaces, and complete low voltage designs all need to be considered.

3.1 Power Supply Design

When designing the power supply for a low voltage system, the characteristics of that system must be taken into account. A supply must exist for each operating voltage on a board. This could mean a number of different voltages for logic, communications, displays, and other functions. For each voltage generated, the designer must determine the worst-case current consumption of the system. Number of supplies and their capacity determine the power supply design.

3.1.1 BATTERIES

A major application of low voltage devices in embedded systems is in portable systems that require battery operation. The type of battery the designer chooses affects the power supply design. Batteries are characterized by two general discharge curves. Alkaline batteries have a constantly decreasing voltage output (Figure 3-1a). Ni-Cad batteries (as well as Nickel-Hydride, Lithium and others) have a relatively constant voltage through the majority of their life (Figure 3-1b). Internal chemical reactions and the number of cells inside each battery determine the output voltage.

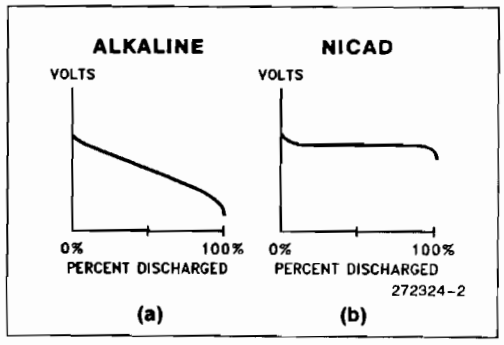

Figure 3-1. Typical Battery Discharge Curves

The battery discharge curve and voltage requirements of the system components determine whether voltage regulation is required. A system utilizing components with narrow voltage specifications (3.3V ±10%, for example) using Alkaline batteries, would require a voltage regulator to translate the battery output voltage to something usable by the system. NiCad batteries, on the other hand, provide a stable enough voltage over their lifetime that the designer could avoid using a regulator.

Although the constant voltage type of batteries may appear to be the most efficient choice, a number of other tradeoffs exist such as overall capacity, rechargeability, temperature sensitivity, safety, volume, weight and cost. When deciding on a battery, these characteristics, or combinations of these characteristics should be considered. For example, an application may require maximum capacity, but minimum weight, so capacity per unit weight may be the most important combination. Also, rechargeability may or may not be important to the design. Secondary batteries, such as NiCad may be recharged, while primary batteries such as alkaline may not.

3.1.2 VOLTAGE REGULATION

Voltage regulators may be required in systems where battery output voltage and system V_{CC} requirements do not match. For example, three Alkaline AA batteries would discharge from about 4.5V to about 2.7V over their lifetime. A device specified for operation from 3V to 3.6V would require this supply to be regulated. Voltage regulated systems have the advantage of having a stable, constant V_{CC} for the life of the batteries. The disadvantages include extra board space needed for the regulator(s) and decreased battery life due to regulator efficiency losses. Most regulator efficiencies run from the mid-80 percent to low-90 percent range.

In low power designs, voltage regulation is usually performed with pulse-skipping regulators, current-mode pulse width modulators (PWM's) or voltage-mode PWM's. Each of these methods has certain advantages and disadvantages. Important regulator characteristics include output ripple amplitude and frequency, efficiency, quiescent current, transient response and physical size requirements. Individual designs will have different requirements, and these issues should be considered when designing the system power supply.

3.1.3 CALCULATING SYSTEM POWER CONSUMPTION

When designing a system power supply, it is important to understand system current requirements. Knowing the worst-case current consumption allows the designer to choose the voltage regulation method, estimate system battery life and determine system packaging requirements (heat dissipation, etc.).

Many factors determine total system current. Device current, discrete component current, and leakage current are all factors. The general formula for current consumption, I is:

Equation 3.1: $I = V \bullet C \bullet f = C \bullet dV/dt$

Where: **V** = voltage, **C** = capacitance, and **f** = frequency. This formula must be used for all components in the system to estimate current consumption. The term "estimate" is used because, until the design is completed, fabricated, and tested, there is no way to know the exact current consumption. Actual current will depend on the way the software controls the hardware in the system, temperature and other factors. To estimate current consumption, some assumptions are required to keep the calculations from becoming too detailed. The various contributors to system current consumption are detailed in the following sections.

3.1.3.1 Device Current

There are two parts to device current consumption: core current and I/O current. Core current is typically a large percentage of the total device current. I/O current is a smaller percentage, but still a major contributor to total system current.

Core current is the current consumed by internal transistors switching and the charging and discharging of internal capacitances. In this situation, V in Equation 3.1 is equal to V_{CC}. All of the internal nodes switch between V_{CC} and ground. The frequency term is the device operating frequency. The Capacitance term, slightly more complicated, is an equivalent capacitance accounting for all internal node and trace capacitances. To determine this value, I_{CC} is measured at a known frequency and voltage with output pins inactive, Equation 3.2 is then solved for the equivalent capacitance (Equation 3.3). This value applies to all operating frequencies.

Equation 3.2: $I_{core} = V_{CC} \bullet C_{eq} \bullet f$

Equation 3.3: $C_{eq} = I_{core}/(V_{CC} \bullet f)$

I/O current is the current consumed by outputs switching. To determine I/O current, V in Equation 1 is the voltage swing of the output, V_{CC} in a normal, CMOS system. The capacitance term is the sum of all input capacitances of devices connected to the output plus the capacitance of the PCB trace (approximately 2 pF/in.). The frequency term is the switching frequency of the output, not necessarily the operating frequency of the system. The frequency term can be determined by approximating how often an output switches during a particular time period. The measure of time in an embedded system is typically a processor bus cycle. Using the 80C186 embedded processor as an example, the major contributors to I/O current are the Address/Data bus, RD#, WR# and ALE outputs. Depending on the application, other outputs may switch frequently enough to be included in current calculations. These signals are used as a simple example.

Because of the multiplexed Address/Data bus on the 186 device, data is written and read on the same pins that drive address information. Thus, read and write cycles need to be treated separately. An analysis of bus cycles for a typical application indicates that about 80 percent are read cycles (data reads and instruction fetches) and the remaining 20 percent are write cycles. Equation 3.4 represents the Address/Data bus current for a write cycle.

Equation 3.4: $I_{write} = (V \bullet C \bullet f) \bullet (1/n) \bullet (x + y)$

where:

- V = V_{CC}
- C = Load capacitance per pin
- f = Processor operating frequency
- n = Number of clocks per bus cycle (4 at zero wait states)
- x = Number of pins switching during the address phase
- y = Number of pins switching during the data phase

Because of the random nature of data being read/written, it is a reasonable assumption that approximately half of the Address/Data pins will switch during each phase of the bus cycle (address and data). For a read cycle, the equation will be identical, with y = 0, as shown in Equation 3.5.

Equation 3.5: $I_{read} = (V \bullet C \bullet f) \bullet (1/n) \bullet (x)$

The control signals: RD#, WR# and ALE are driven during each bus cycle. Equation 3.6 shows the current for the combined control signals. Each of these signals switch twice during a bus cycle, once to go active and once to go inactive or vice-versa.

Equation 3.6: $I_{control} = (V \bullet C \bullet f) \bullet (1/n) \bullet (x)$

- V = V_{CC}
- C = Load capacitance on the control signal pin
- f = Processor operating frequency
- n = Number of clocks per bus cycle (4 at zero wait states)
- x = Number of times the control signal switches during a bus cycle
  - ALE: x = 2
  - RD#/WR#: x = 2

Total I/O current, Equation 3.7, is calculated by adding equations 3.4, 3.5 and 3.6.

Equation 3.7: I/O current = $I_{read} \bullet (0.8) + I_{write} \bullet (0.2) + I_{control}$

The 0.8 and 0.2 factors come from the percentage of read and write bus cycles in a typical system. A similar analysis is required for all devices within a system. The system processor will usually be the most difficult. Once this is done, many of the approximations made apply to calculations for the rest of the system.

3.1.3.2 Discrete Component Current

Current consumed by discrete components must also be factored into total system current. Current consumed by pull-up or pull-down resistors, voltage dividers, transistors or other discrete components must be included. These are relatively straightforward calculations, and do not require elaboration.

3.1.3.3 Leakage Current

Leakage currents from devices are typically very small and can usually be ignored when calculating system current consumption. When a CMOS device is in a static condition, all p and n transistors inside the device should have their gates driven to V_{CC} or Ground, turning them on or off. The equivalent resistance of an "off" transistor is approximately 5 MΩ, the equivalent resistance of an "on" transistor is typically less than 100Ω. Some current flows through this resistive path to ground. Typically, this amount will be in the μA range.

Leakage current can become more significant. If a device is running at a given V_{CC}, there may be situations where an input is not driven close to V_{CC}. The p and n transistors of the device input buffers will not be com-

pletely turned on or off. This lowers the equivalent resistance through the transistors to ground. On some devices (AC logic, as an example), at $V_{IN} = V_{CC} - 2.1V$, the device may consume as much as 1.5 mA additional current **per input pin**. Typically, this value is from 100 μA to 200 μA per input, but still measurable. This additional current exists unless the input to the device is driven close to V_{CC}. Most logic data sheets list this specification as ΔI_{CC}, given as a maximum current per input pin with the input at $V_{CC} - 2.1V$.

3.1.4 SUMMARY: CALCULATING POWER CONSUMPTION

For a complete system design, total system current consumption must be calculated. This analysis determines power supply design, product packaging and battery life. Although an exact calculation may not be possible (or practical), a worst-case analysis with reasonable approximations allows a reliable system design.

3.2 Mixed Voltage System Design

One of the difficulties in designing a low voltage system is product availability. There are a large number of manufacturers producing low voltage products (3V or lower). Embedded processors, like the Intel 80L186EA, EB and EC have been available at 3V for some time. Memory and logic are also widely available. Some peripheral components are still not available at low voltages. Many of these components require redesign to operate at low voltages. Until these components are replaced with low voltage versions, systems requiring their functionality have to utilize the 5V version. This causes tremendous headaches for system designers. System cost, complexity and power consumption all suffer due to the lack of a complete selection of low voltage devices.

3.2.1 INTERFACING 3V AND 5V COMPONENTS

Somewhere within a mixed voltage design, the designer must interface 3V and 5V devices. This can happen in 3V to 5V translation, 5V to 3V translation, bidirectional translation or 3V and 5V devices residing on the same system bus. As simple as these interfaces may seem, there are still problems to be solved.

3.2.1.1 3V to 5V Interface

One way to translate a 3V output to a 5V CMOS level is with ACT/HCT logic. These parts accept a TTL level input and give a CMOS output. When operating at 5V, these parts see a 3V level input the same as they would see a 5V TTL level input. The output will be a 5V CMOS level. This is a relatively straightforward approach, but it does have the problem of additional current consumption (ΔI_{CC}).

The JEDEC standard for 3.3V outputs is compatible with TTL input requirements. If the 5V device connected to the 3V output already has TTL compatible inputs, no interface logic is required.

3.2.1.2 5V to 3V Interface

Unfortunately, a 5V output cannot be connected directly to a 3V input. Maximum operating conditions for most devices specify the maximum voltage on any input pin. This number is typically V_{CC} + 0.5V (LVTTL and unregulated LVCMOS), and on some 3V devices is V_{CC} + 1.0V (regulated LVCMOS). A 5V CMOS output high typically drives close to V_{CC}. When the maximum input voltage is exceeded, the ESD (Electrostatic Discharge) protection diode on the input of the device forward biases and current flows into the 3V V_{CC} (Figure 3-2). This causes unacceptably high current consumption. Connecting a 5V output to a 3V input leads to reliability problems. Device manufacturers are developing input buffers for low voltage devices that are 5V tolerant, but until these devices are widely available, other solutions are required.

A 5V bipolar TTL device may be connected directly to a 3V input. Some newer 5V CMOS devices have outputs with reduced voltage swings. These devices have 3V compatible outputs and may be connected directly to 3V inputs. The designer must be careful when regulating the 5V and 3V supplies. If the 5V supply goes to 5.5V and the 3.3V supply goes to 3.0V, the 3V input specification could be violated.

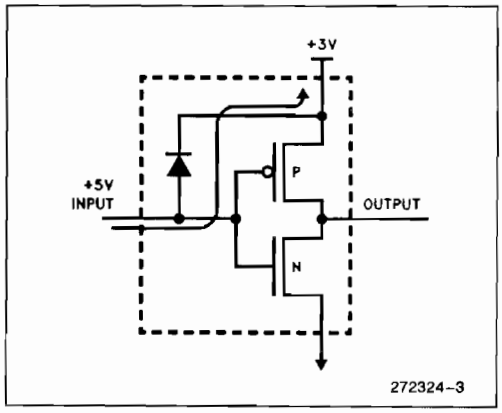

Figure 3-2. 5V Output Driving a 3V Input

There is a way to work around the problem of driving a 3V input with a 5V output. Series resistors can be placed on the outputs of the 5V device. The goal is to drop the voltage to a level acceptable by the 3V device. The value of the resistor must be chosen to limit the

current into the 3V device. The major sacrifice in this solution is speed. If the system can handle the speed degradation, then it is a simple, inexpensive way to translate from 5V to 3V. As a final point on this subject, if this solution is implemented, considerations must be made for system power-up. If the 5V supply ramps much faster than the 3V supply, the substrate diode may still be forward biased temporarily, leading to reliability problems. This situation must be taken into account when determining the resistor value. The designer must either assume the worst case where the 5V part V_{CC} = 5.5V and the 3V part has V_{CC} = 0V or the 3V and 5V supplies must be sequenced.

3.2.1.3 Voltage Translation with Open Drain Outputs

Open drain output devices provide a simple way to convert from 3V to 5V and vice-versa (Figure 3-3). All that is required is an external pullup resistor to the desired output voltage. If the open drain device outputs a logic "1", there is virtually no current consumption penalty for the conversion. If the output is a logic "0", there is a current path to ground, but a high resistance pullup will limit the amount of current (at the cost of speed).

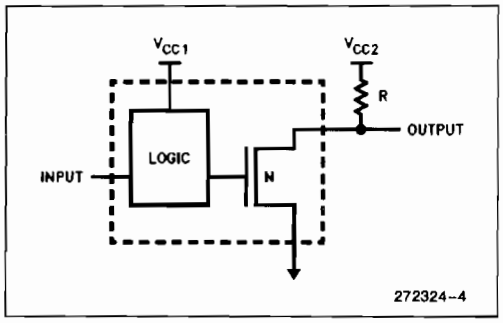

Figure 3-3. Voltage Translation with Open Drain Outputs

If an open drain device is not available, the function can be easily duplicated using an external MOSFET and a resistor. This circuit will be identical to Figure 3-3, except the transistor will be external to the device. The output to be translated connects to the gate of the transistor. An n-transistor connected to ground with a pullup to V_{CC} will act as an inverter for the output. A p-transistor connected to ground with a pullup resistor will not invert the output. The switching time required by the external MOSFET must be considered when determining system timing.

3.2.1.4 Bidirectional Translation

The methods described above all work well for translating unidirectional signals. What happens if a 5V peripheral resides on a 3V bus? Ideally, a designer could place a 5V device and a 3V device on the same system bus. Unfortunately, a floated 3V output will be damaged when a 5V part drives the bus. Depending on the state of the inputs to the 3V device output buffers, it is possible that the p-transistor will turn on (Figure 3-4). If this device turns on, the 5V supply and 3V supply are shorted together. Even if this situation does not occur, the 5V signal still forward biases the ESD protection diodes inside the 3V device, creating a situation similar to a 5V output driving a 3V input. Fortunately, a number of manufacturers are producing buffers with two V_{CC} pins. The devices translate bidirectionally between the two V_{CC} values. These are the most practical solution for bidirectional translation between 3V and 5V. System busses must operate at a single voltage with other voltages buffered.

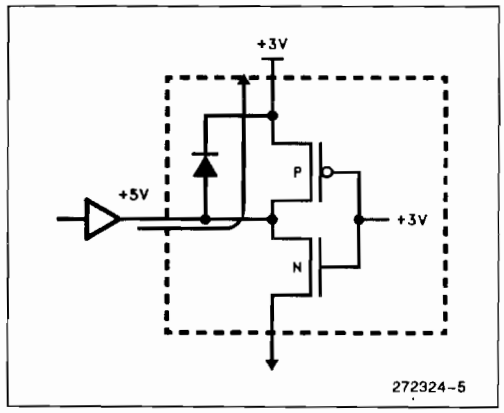

Figure 3-4. 5V Device Driving a Floated 3V Output

3.2.2 DISADVANTAGES OF MIXED VOLTAGE SYSTEMS

It is obvious at this point that there really are no advantages to mixed voltage systems over complete 3V designs. They only exist because of the absence of a complete selection of low voltage devices. During the industry's transition to 3V, the objective of designers is to minimize the disadvantages of mixed voltage systems. Two of the drawbacks of mixed voltage systems are voltage supply requirements and additional current consumption.

3.2.2.1 Multiple Supplies in Mixed Voltage Systems

A major disadvantage of mixed voltage systems is the requirement of multiple voltage supplies. A typical system may require 3V (major components, memory), 5V (older peripherals, small displays), ±12V (RS-232 communications) and even higher voltages (backlit LCD displays, etc.). One goal in designing a mixed voltage system is to minimize the number of required voltages and the number of devices used to create them. To avoid the extra chip-count associated with creating multiple voltages, some manufacturers, offer one-chip solutions to provide multiple voltages from 2 or 3 cells. The designer can also take advantage of parts with internal charge pumps that take 3V or 5V inputs and create the output voltage levels they require. Although multiple voltages can easily be created with a minimal number of chips, the designer still pays the price for having non-3V parts in the system: battery life. As technology moves forward and 3V designs gain momentum, more components will be available at low voltage (3V or less). This eliminates the requirement for multiple system voltages and the added system cost and complexity associated with creating them.

3.2.2.2 Additional Current Consumption in Mixed Voltage Systems

Interfacing 3V and 5V devices in a mixed voltage system is unavoidable. Regardless of how a designer implements these interfaces, they will all have one common characteristic, additional current consumption.

Some devices have an additional specification called ΔI_{CC} (or I_{CCT}). This specification defines the additional current consumed, per input pin, if an input high voltage is at $V_{CC} - 2.1V$. This situation closely resembles using ACT or HCT logic for 3V to 5V translation. This number can be up to 1.5 mA per input pin. Consider a unidirectional, 16-bit bus translated from 3V to 5V using ACT logic. In a worst-case situation, this can be a major source of continuous current consumption. This is a maximum value, typically the extra current amounts to 100 μA to 200 μA per input pin. Additionally, this specification only applies to a logic "1" input, and typically, only a fraction of the 3V inputs are high at any time.

The reason behind the extra current consumption when using ACT/HCT logic for 3V to 5V translation lies in the input buffers (Figure 3-5). If the input of the device is driven all the way to V_{CC}, the p-transistor turns completely off and the n-transistor turns completely on. This can be roughly modeled as 5V connected to ground through a 5 MΩ resistor. As shown by the graph in Figure 3-5, the only current flowing is leakage current, almost nothing.

As the voltage on the input moves farther away from V_{CC}, the input transistors move closer to their saturation region. The resistance path through the transistors to ground decreases from the initial 5 MΩ. This increases the current flow through them. The graph shows this current to be on the order or 150 uA per input with a 3V input (at room temperature).

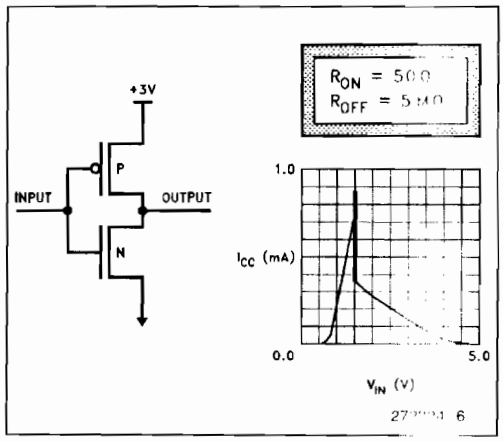

Figure 3-5. Current When Using ACT/HC[illegible]ic for Voltage Conversion

This leads to a valid question. If there is such a penalty for creating mixed voltage systems, is the system better off running at 5V? An analysis of system current consumption must be done for the pure 5V and mixed voltage cases. If only a small part of the system can operate at 3V, the extra power for voltage translation [illegible] offset the benefit of using 3V parts. This accentuates the need to have a complete selection of devices that operate at 3V.

3.2.3 SUMMARY: MIXED VOLTAGE SYSTEMS

There are a number of considerations a designer must make when designing a mixed voltage system. Interfacing 3V and 5V logic must be done carefully. There are many solutions to do this, but if done incorrectly, the system eventually fails. Many manufacturers provide simple solutions to do unidirectional and bidirectional transfers. These solutions are probably the easiest to implement and consume a minimal amount of power.

A mixed voltage system, by definition, consumes more current than a pure 3V system. This added current consumption comes from different sources. Any method used to translate from one voltage to another requires additional current. A mixed voltage system also has more devices than a pure 3V system. Extra devices are needed for voltage translation and creating required system voltages. In addition to drawing current, these extra devices increase system size, cost and complexity.

As technology advances and manufacturers redesign current parts, complete systems will be able to operate at 3V. Until that time, mixed voltage systems must exist. Although a mixed voltage system requires more power than an entirely 3V version, it still consumes less power than a 5V version (if properly designed). Mixed voltage systems will exist in some form for a long time. Right now, they exist because of the conversion from 5V devices to 3V devices. They will continue to appear as industry makes the steps to even lower operating voltages. Some devices already operate at 2V and below. Although this information applies specifically to 3V/5V systems, the concepts apply to any mixed voltage system.

3.3 Single Voltage System Design

Complete 3V designs have few disadvantages, the exceptions being operating speed and noise immunity. Advantages include: longer battery life, less heat and lower noise emissions. The advantages of a complete 3V system are significant compared to the minimal design effort required to work around the disadvantages.

3.3.1 DEVICES DESIGNED FOR LOW VOLTAGE OPERATION

Currently, some devices are designed for 3V only operation. Devices optimized for 3V operation are generally specified for operation from 3.0V to 3.6V. These devices exhibit the performance of their 5V counterparts at significantly lower power. The high performance and narrow operating voltages of these parts [illegible] their benefit to embedded, battery operated designs.

If a fabrication process is optimized for 3V operation, gate oxides are thinner. The increased gate capacitance allows a 3V device to function at the same speed as a 5V device produced on a non-3V processes. Oxide breakdown is the major drawback to optimizing a process for 3V. A thinner gate oxide implies a lower oxide breakdown voltage. Devices produced on the optimized process will not be able to operate at 5V without damaging the part.

3.3.2 DEVICES DERATED FOR LOW VOLTAGE OPERATION

Many low voltage products are derated versions of their 5V versions. These devices are specified, typically, for operation from 2.7V to 5.5V. The only sacrifice of running a product designed for 5V at 3V is speed. Device speed is a function of internal switching speeds. Switching speeds are proportional to device current, a function of gate oxide capacitance and gate voltage within a device. A device produced on a process not optimized for 3V will have a thicker gate oxide, decreased gate capacitance, reducing current The combination of lower gate voltage and capacitance limits operating speeds for derated 5V devices. For many embedded systems, speed is not critical. Using devices that operate from 2.7V to 5.5V allows the designer to use inexpensive alkaline batteries and avoid the added device cost and efficiency losses of voltage regulators.

3.3.3 NOISE GENERATION BY LOW VOLTAGE DEVICES

Overshoot, undershoot and ground bounce all relate to dV/dt, dV being the output voltage swing and dt being the output switching time. In low voltage devices, obviously, voltage levels decrease, lowering the dV term. This applies to all low voltage devices. Devices designed specifically for low voltage operation, which have smaller geometries and thinner gate oxides, have switching times comparable to the 5V versions. Therefore, the dt term remains relatively constant. Low voltage devices which are derated versions of 5V parts have slower switching times when operating voltage decreases. The dt term increases for devices derated for low voltage operation. Derated devices generate very little noise.

Regardless of the exact amount of noise a low voltage device generates (derated or not), any low voltage device will produce less noise than its 5V counterpart. dV/dt and practicing good PCB design techniques (multi-layer, bypass capacitors, etc.), should eliminate noise as a problem.

3.3.4 NOISE MARGINS IN LOW VOLTAGE SYSTEMS

In devices with CMOS compatible inputs, input high (V_{IH}) and input low (V_{IL}) voltage specifications are a function of operating voltage. The difference between V_{IH} and V_{IL} defines the noise margin on the input (Equation 3.8).

Equation 3.8: Noise Margin = $V_{IH} - V_{IL}$

Typically, V_{IH} is $0.7*V_{CC}$ and V_{IL} is $0.3*V_{CC}$ on CMOS compatible inputs. This defines a noise margin of $0.4*V_{CC}$. If operating voltage decreases from 5.0V, noise margins decrease proportionally.

Devices with TTL compatible inputs will not see any change in noise margins when operating at 3V. TTL inputs define V_{IH} as 2.0V and V_{IL} as 0.8V. The standards for input levels on 3.3V devices are defined to be compatible with the TTL standard. This is why a 3.3V output can directly interface with a 5V device with TTL inputs.

Noise susceptibility of low voltage systems is not as significant as it may appear. Low voltage systems generate less noise than 5V systems. If a system generates less noise, decreased noise margins on device inputs become less important. Additionally, good Printed Circuit Board layout techniques should eliminate major noise issues.

4.0 SYSTEM POWER MANAGEMENT

When creating a battery operated system, the designer is concerned with extending battery life as long as possible. Reducing operating voltages is a relatively simple way to achieve a significant reduction in power consumption. Voltage reduction is an excellent start to increasing battery life, but numerous power management techniques exist to provide even lower power consumption.

4.1 Device Power Management

Many Intel embedded processors have modes of operation designed to reduce current consumption. Many static-design embedded processors have Powerdown mode, disabling the clock input to the device. Disabling the clock input eliminates transistor switching within the device, reducing current consumption to leakage current (microamps). Other devices have Idle mode which disables the clock to the CPU, but keeps integrated peripherals active. This decreases current consumption by a smaller amount than Powerdown mode, but is excellent for devices requiring peripherals remain active at all times. Finally, many embedded processors offer Powersave mode. Powersave mode internally divides the clock input to the device. Because processor current consumption is approximately a linear function of clock frequency, Powersave mode significantly reduces current consumption during execution of non-critical sections of code. Including one or more of these power management functions on silicon reduces board space required to create them separately, if they can be created at all (Idle mode, for example).

Other devices that commonly appear in embedded systems also offer power management features in addition. Some Intel Flash memory devices offer a powerdown mode, reducing current to leakage levels. Other memories enter a "sleep" mode when not being accessed, reducing current consumption. Even some signals allow their output to be shut down to decrease current consumption.

Many devices exist offering power management features which enhance low voltage operation. When system designers select components for their systems, parts should be chosen which offer benefits best suited for low voltage operation.

4.2 Software Power Management

Power management schemes need to be implemented in software as well as hardware. Software controls power management modes on Intel embedded processors and can be used to control power management on external devices (using port pins, etc.). Software can also be implemented to reduce switching on address and data busses and processor outputs.

4.2.1 IMPLEMENTING DELAYS

Many applications require time delays. For example, there may be a minimum time required between successive writes to an LCD display or other peripherals. These delays can be implemented either in hardware (controlled by software) or software. When delay loops are implemented in software, every iteration of that loop will execute a given number of bus cycles. Every bus cycle will drive address information onto the bus and all of the relevant control signals that define a bus cycle and turn on the memory device(s) containing program code. With enough iterations of the loop, a large amount of current is consumed. One solution is to divide the processor clock during these delays. Dividing the clock by a factor of x reduces the number of bus cycles required by x, for the same delay.

Implementing a delay with hardware and software is simple. Timers are used to create the delay. A timer can be programmed to cause an interrupt after a set delay. The timer needs to be programmed with the correct count and enabled. If the delay is long enough, it may even make sense to put the processor into a power management mode (such as Idle mode or Powersave mode). This entire process requires a small number of bus cycles. If the overhead associated with interrupt servicing is a limitation, some processors allow the designer to implement delays without interrupts.

The 80C186, for example, has an input called TEST#. The TEST# pin is sampled during the WAIT instruction. The processor stops execution during a WAIT until the TEST# pin is sampled low. If the output of a timer is connected to the TEST# input, the timer times-out and pulls its output low, causing the processor to continue execution. This solution requires only enough bus cycles to program the timer.

4.2.2 CODE OPTIMIZATION

Program code can also be optimized for minimum power consumption. To minimize switching on the address bus, code should be written to reduce the number of jumps that occur. Interrupt driven code is also preferable to code which depends on polling. The processor can issue a HLT instruction, waiting for an interrupt rather than constantly looping and polling a register bit. The reduction in power consumption using these methods may be slight but still increases battery life.

Power management is very important to low voltage embedded designs. The effort to add power management is minimal but the benefits are significant. Using hardware and software to minimize power consumption increases battery life with little or no increase in system cost or complexity.